Responsible banking in the age of AI

Why the introduction of the **FCA's Consumer Duty** means firms need to get serious about the responsible use of AI.

sas

Highlights

Consumer Duty requires firms to act in good faith, avoid causing foreseeable harm, and support retail customers in pursuing their financial objectives.

Artificial Intelligence (AI) can help firms achieve these goals—but can also pose significant compliance risks if models are inaccurate, outdated, or improperly governed.

Model risk management and model lifecycle governance must be effective to ensure AI is used responsibly and provides good outcomes for consumers.

SAS can help firms establish comprehensive governance frameworks for AI models, with integrated tools that can accelerate model development, validation and overall governance.

Rapid deployment with SAS will help firms meet the FCA's challenging Consumer Duty implementation deadlines and ongoing compliance, avoiding the significant cost of remediating models manually.

What is the problem?

After more than half a century of false dawns, Artificial Intelligence has finally crossed the chasm to mainstream adoption. Even in banking's highly regulated environment, the potential benefits of AI have become too significant to overlook. Banks, building societies and regulators must now work together to ensure this powerful technology is used responsibly.

Firms have been quick to realise that AI models aren't just useful in traditional domains such as credit risk and stress testing–they can also be used to aid decision-making and automate processes in areas such as marketing, customer experience, and collections. As the number and complexity of AI models in these areas increases, these teams need to ensure their governance and model risk management are suitably robust.

This poses risks. AI models are complex. It's challenging to make them transparent and explainable. It's easy to inadvertently build inaccurate or biased models, and the effects can range from reputational damage to regulatory fines. High-profile controversies such as the Apple Card sexism scandal have highlighted the danger of models discriminating against certain groups, even when the model builder has been careful to exclude protected characteristics such as age, sex, race and religion from their training datasets.

Inaccurate or biased models are a serious business liability. In the short term, if your models prevent you from lending to a customer who can afford the loan, you're missing out on good business. On the other hand, if you lend to someone who can't repay, they'll end up in collections, or claiming intolerable harm to the ombudsman. In the long term, the risk of mis-selling products can cost hundreds of millions to remediate, as the PPI scandal has shown.

Moreover, as financial services become increasingly commoditised, customer experience is key. Inappropriate use of models is likely to impact customer experience, leading to lower satisfaction, reduced loyalty, increased complaints, and higher costs.

What the regulators are doing

The FCA's new Consumer Duty rules (PS22/9) place a wide-ranging set of obligations on firms to rethink every aspect of their relationships with retail customers—including their use of AI in decision-making. Biased AI models are explicitly called out as contravening the Consumer Duty Cross-Cutting Rule on acting in good faith:

> "Examples of not acting in good faith [would include using] algorithms, including machine learning or artificial intelligence, within products or services in ways that could lead to consumer harm. This might apply where algorithms embed or amplify bias and lead to outcomes that are systematically worse for some groups of customers, unless differences in outcome can be justified objectively."

FCA, Finalised Guidance (FG22/5)[i]

[i] https://www.fca.org.uk/publication/finalised-guidance/fg22-5.pdf

Consumer Duty also requires firms to avoid causing foreseeable harm to customers. One of the key requirements is to ensure that products and services are designed to fit the needs of their target markets and represent fair value. Again, this has consequences for the use of models in product design and pricing, customer segmentation and marketing. If an inaccurate model causes a firm to sell inappropriate products to potentially vulnerable groups, the penalties could be significant.

The key word is "foreseeable", and there will doubtless be debate about what types of harm can be foreseen. Yet even if the regulations are open to interpretation, firms will need strong governance in place to demonstrate that they have taken all reasonable precautions to keep their models free from bias from inception and throughout their lifecycle. Depending on their appetite for operational risk, firms may seek statistical proof that their decisions are fair.

Consumer Duty and AI: risk versus reward

AI benefits

AI can help firms provide better access to financial services for consumers with non-standard histories

AI can help firms tailor products to individual consumers to help them achieve their financial objectives

AI can help firms target products, information and support to appropriate customer groups, helping them make better financial decisions

Consumer Duty compliance risks

AI could widen existing gaps in society and lead to firms inadvertently exploiting vulnerable customers

AI could implement harmful price discrimination and adopt behavioural biases that negatively impact customer outcomes

AI could market inappropriate products or provide misleading information to customers, raising the risk of mis-selling

Challenges for compliance

The FCA's deadlines for the initial implementation of the Consumer Duty rules are fast approaching, and there's concern in both banking and regulatory communities that firms will not be ready in time. Even so, the FCA is taking a robust position on enforcement, warning banks and building societies that serious breaches will not be overlooked:

> "Our supervisory and enforcement approach will be proportionate to the harm—or risk of harm—to consumers, with a sharp focus on outcomes. We will prioritise the most serious breaches and act swiftly and assertively where we find evidence of harm or risk of harm to consumers."

Sheldon Mills, FCA Executive of Consumers and Competition[ii]

[ii] https://www.fstech.co.uk/fst/Firms_Who_Ignore_Consumer_Duty_Will_Face_Swift_Action_Warns_FCA.php

One of the biggest challenges is that Consumer Duty has a broad remit. All consumer processes need to be reviewed for compliance. Any use of AI, machine learning or other automated decisioning technology that could impact consumers is therefore in scope.

Rigorous model governance must therefore be extended not just across credit, market, capital risk and stress testing models, but across models used in marketing, customer service, collections and other customer-facing areas too.

The first challenge, then, is to identify all models that require scrutiny. This can be a difficult task even if a firm operates only in the UK and has a relatively simple product set; for larger multinational organisations, it will be a significant undertaking.

Once the firm has taken inventory of its models, the process of reviewing them and determining their effectiveness will consume considerable management time—and keeping these reviews up-to-date and auditable is an even larger ongoing effort.

Key questions for responsible model governance

How can you centralise and standardise the governance of all models used by your firm?

Can you prove to regulators that all your model-assisted decisions are fair to consumers?

How long does it take you to detect, retrain and redeploy an inaccurate or biased model?

How do you manage and report on incidents when a model makes poor decisions?

Can you identify affected customers quickly and respond proactively to avoid intolerable harm?

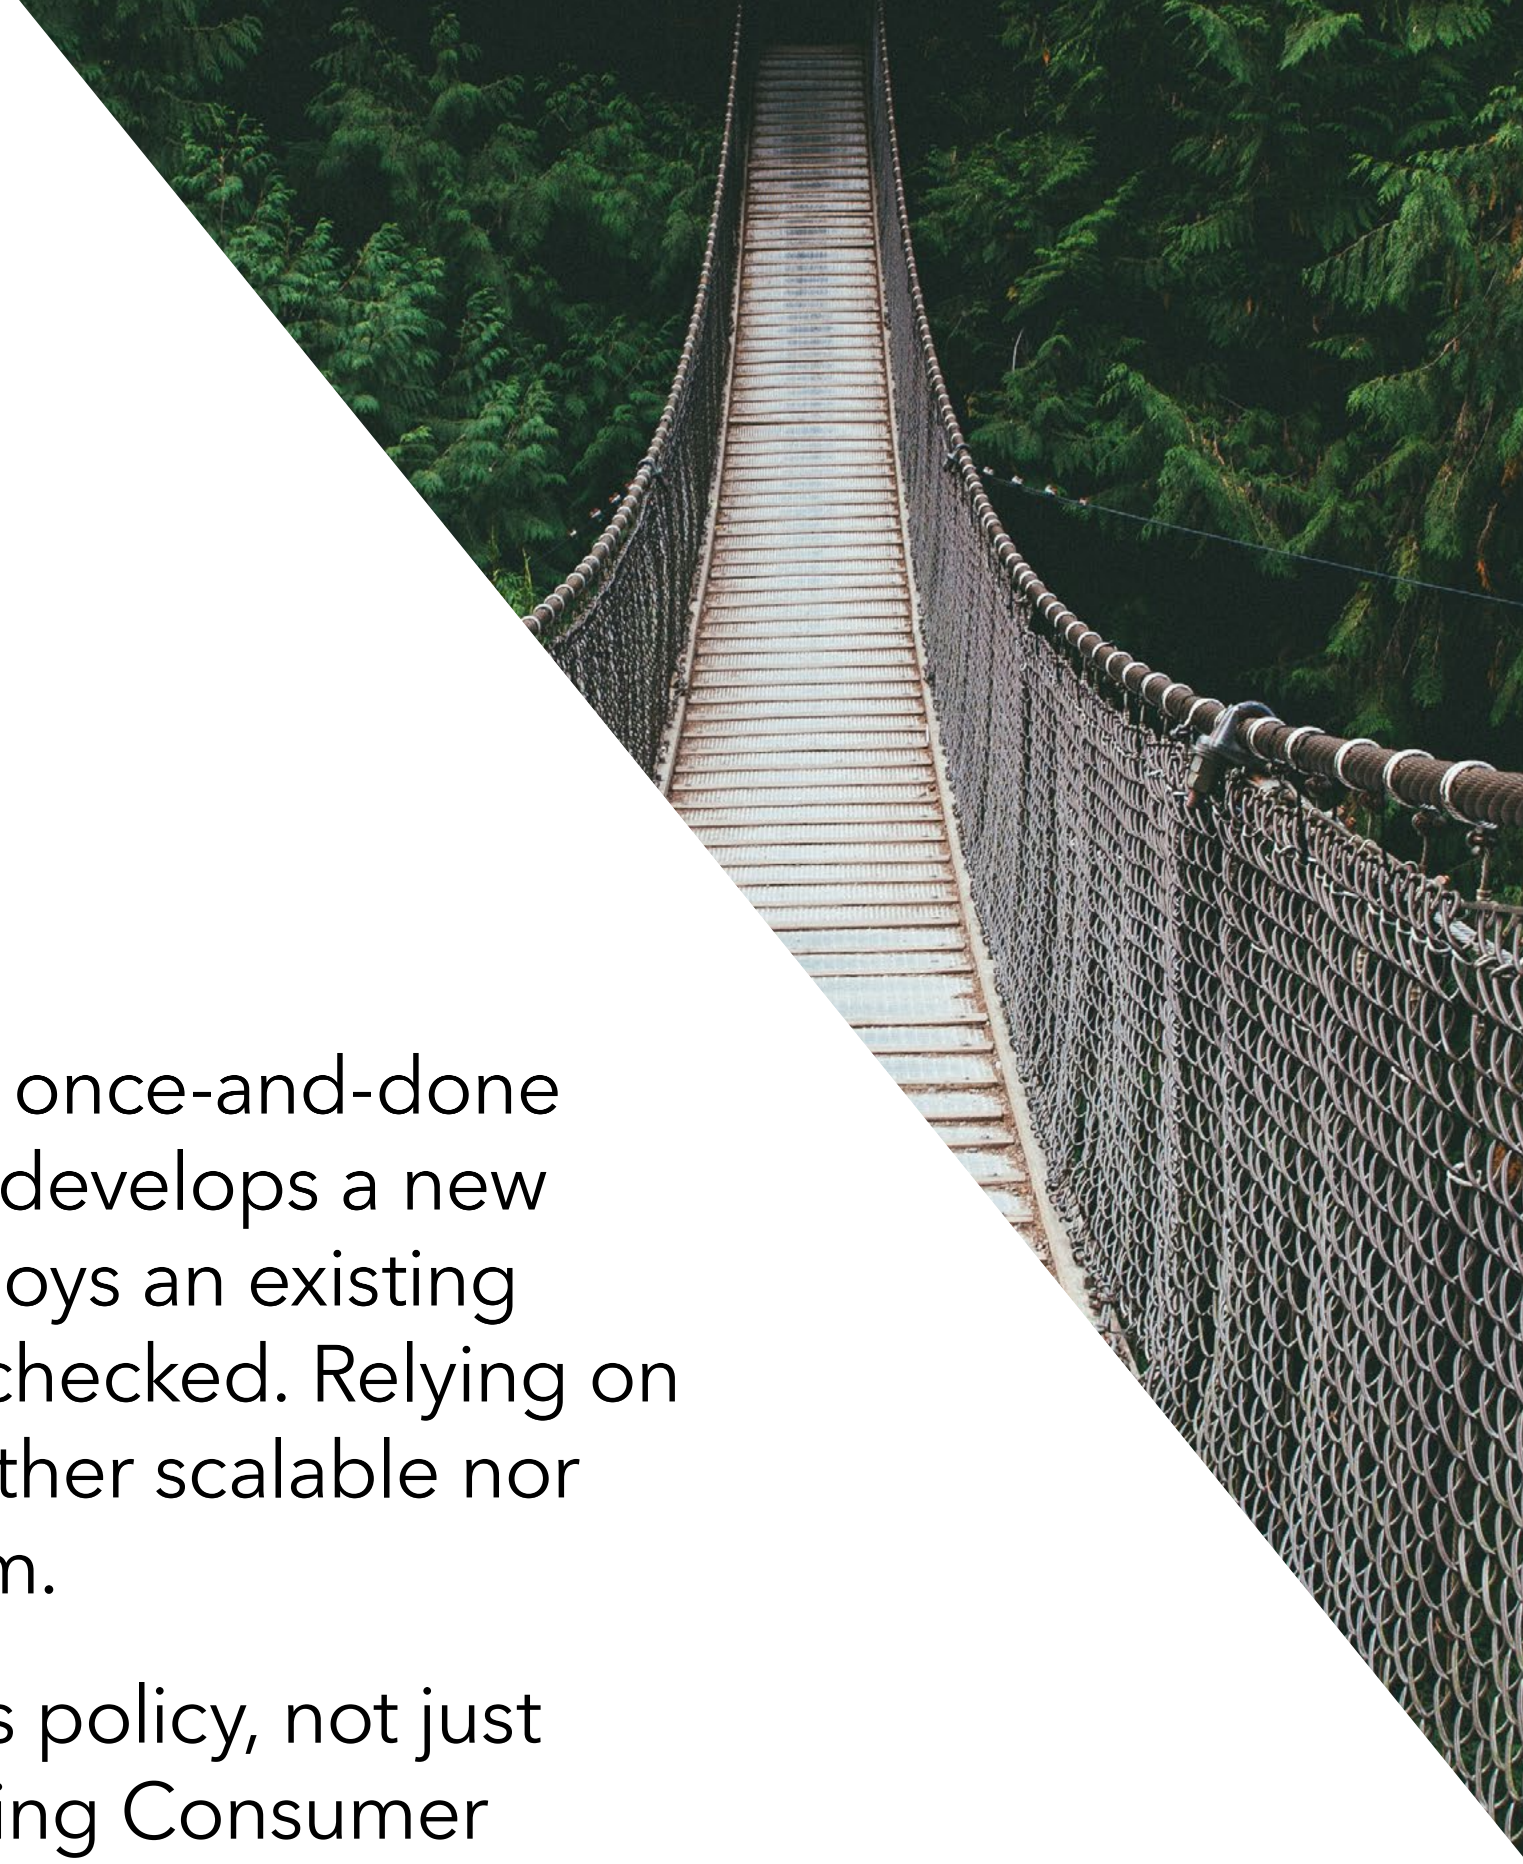

How to address the problems

In the short term, some banks and building societies are proposing to meet the imminent Consumer Duty deadline by taking a bare-minimum approach: a one-off review and remediation exercise to address any issues with accuracy, recency and bias in the models they are currently using in production.

This is an expensive and risky approach. Most firms don't have the resources in-house to complete such a review in time, so they will need to hire a small army of highly skilled contractors at considerable cost.

Moreover, compliance isn't a once-and-done exercise. Every time the firm develops a new model or retrains and redeploys an existing one, they will need to be re-checked. Relying on ad hoc manual reviews is neither scalable nor sustainable in the longer term.

Solving the problem requires policy, not just personnel. Instead of regarding Consumer Duty as a box-ticking exercise, firms should take the opportunity to extend their existing model governance frameworks and tools. By adopting the principles of the Consumer Duty regulation, they can reshape their overarching corporate culture to ensure responsible use of AI both now and in the future.

What if it's too late?

Firms may feel that implementing a comprehensive AI-controlled Consumer Duty framework is impossible within the time available. Even if they narrow the scope to focus on bias detection and mitigation, it could take 6 to 12 months to build the required tooling.

The good news is that there's no need to develop these tools from scratch. SAS provides bias detection and AI ethics tools that work 'out-of-the-box' and can be integrated into a firm's model lifecycle management workflows in a matter of weeks.

What the solution looks like

SAS is already working with several of the UK's largest banks and building societies to implement a solution that can automatically detect and identify ethics issues in AI models from the earliest stages of the model lifecycle.

By highlighting fairness concerns during the model design, development and testing phases, the solution makes it easier to adjust and improve models quickly and cost-effectively. This significantly reduces the risk of inadvertently deploying a problematic model into production—protecting customers from harm and helping to maintain compliance with both the letter and the spirit of Consumer Duty.

By improving observability throughout the model lifecycle management, the solution also enables ongoing remediation. Production models may reveal latent ethics issues at any stage in their lifecycle, so the ability to identify problems with recency, accuracy, and bias both proactively and retroactively is extremely valuable for compliance teams.

Benefits of proactive AI ethics

Sales and marketing

Unlock new revenue opportunities with previously under-served customers

Improve the customer experience to boost loyalty and gain a competitive edge

Compliance and operations

Accelerate and simplify compliance with automated reporting

Reduce headcount costs by avoiding manual model reviews

Risk management

Minimise regulatory penalties by avoiding foreseeable harm to consumers

Eliminate remediation costs by reducing the risk of mis-selling

The road ahead

Determining model effectiveness is a vital first step on the journey to a comprehensive model governance framework that will not only help your financial institution meet the immediate requirements of Consumer Duty, but also give you the flexibility to adapt to evolving consumer expectations, AI developments, and future regulation.

As banking becomes increasingly AI-driven, regulatory focus on model risk management is likely to increase, and consultations such as the Bank of England's SS1/23 indicate the future direction. So, by working with SAS on Consumer Duty today, you are also laying a solid foundation for a well-governed, AI-powered future.

SAS can help your organisation navigate this new world and adopt AI safely, responsibly, and ethically—while unlocking significant benefits in terms of both revenue growth and operational cost reduction.

SAS can help you

1 Prioritise responsible AI adoption

As you integrate AI into your organisation, ensure that it aligns with the ethical principles underpinning the Consumer Duty regulation. SAS can help you establish transparent, accountable, and unbiased AI systems that protect consumer interests.

2 Balance AI benefits and risks

AI can help provide tailored financial services, but it also poses risks of widening existing gaps in society, exploiting vulnerable groups, and enabling discriminatory practices. SAS can help you create a symbiotic relationship between AI and vulnerable consumers to ensure fair treatment for all.

3 Collaborate across disciplines

Firms need to engage multi-disciplinary teams, including data scientists, ethicists, legal experts, and business leaders, to address potential biases, ethical concerns, and regulatory compliance when developing AI models and strategies. SAS can provide the platform for seamless collaboration between all stakeholders across the AI lifecycle.

4 Foster transparency and trust

It's vital to establish clear communication channels with customers, employees, and regulators, and maintain transparency regarding your AI processes and decision-making. SAS can deliver powerful, automated reporting and analytics to build trust in your AI systems and demonstrate your commitment to customer-centricity.

5 Continuously monitor and refine AI systems

Regularly assessing the performance and impact of your AI tools is critical to ensure that they deliver positive outcomes for customers. SAS software and expertise can help you identify and address potential biases or harmful consequences that may arise from AI-driven decision-making.

6 Invest in AI education and training

In an AI-centric world, your C-level executives and employees will need the knowledge and skills to navigate the ethical implications of model-driven decision-making and ensure it stands up to increasing regulatory scrutiny. SAS can help you foster a culture of responsible AI adoption at every level of your organisation.

Why SAS

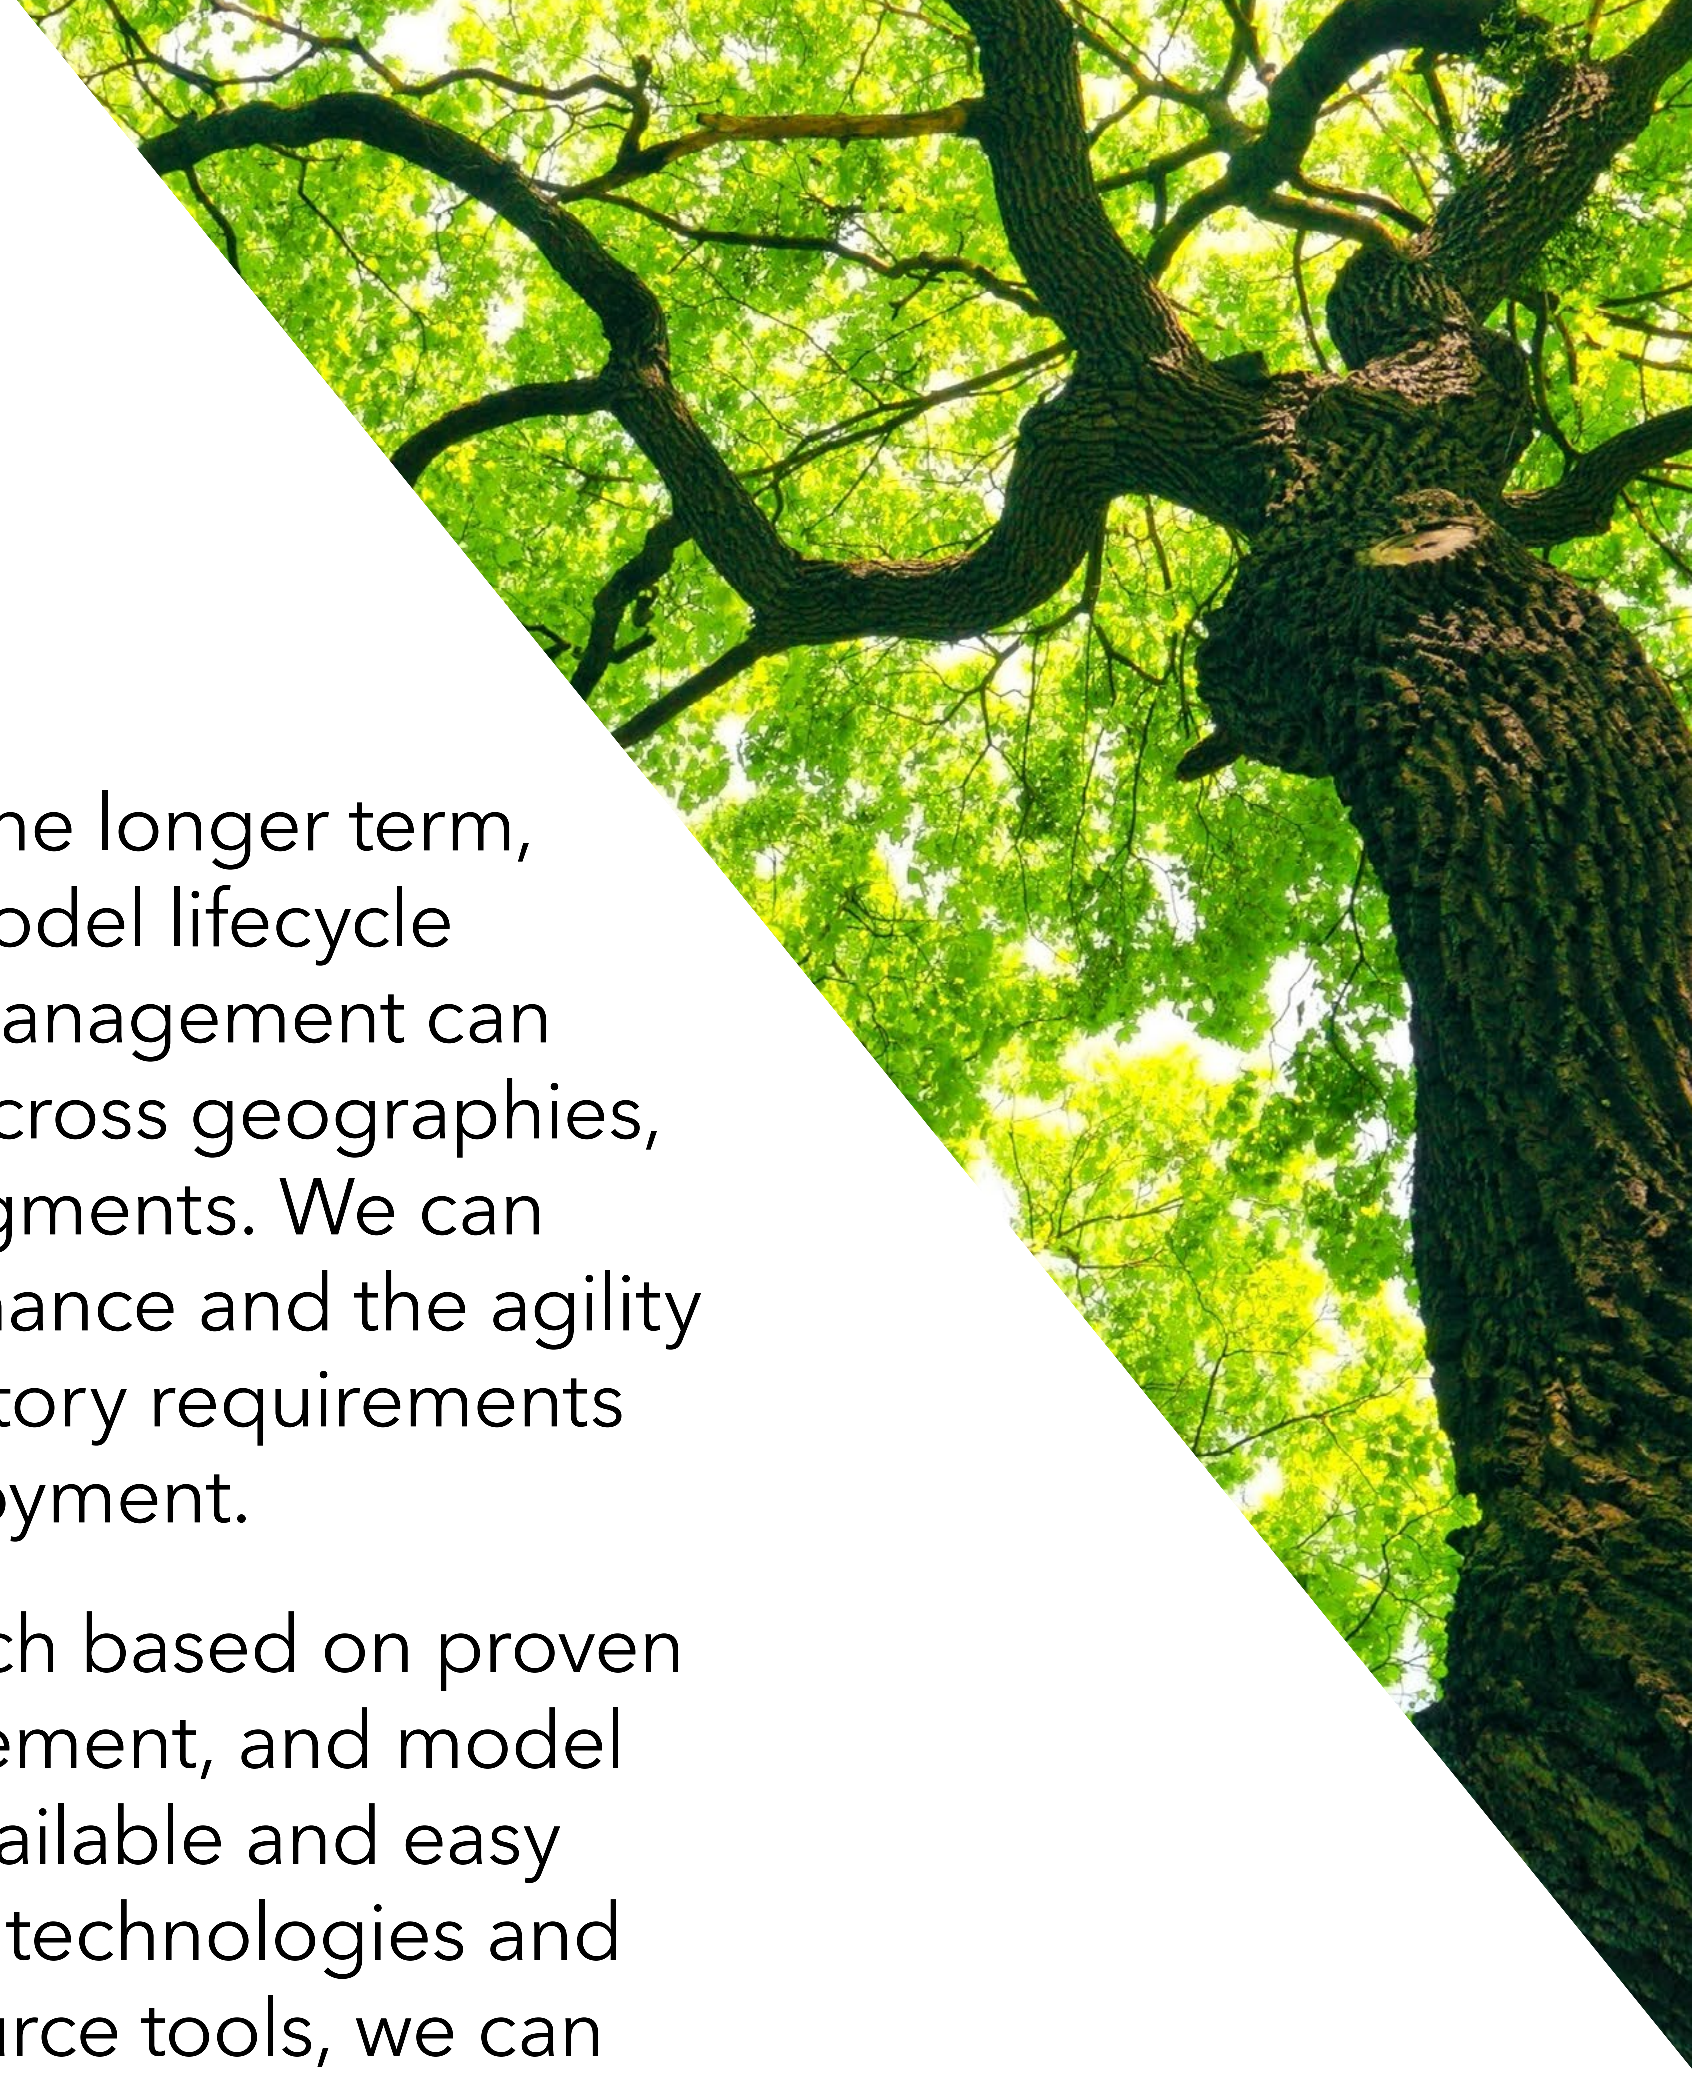

SAS works with over 90% of the top 100 global banks, including 12 Global Systematically Important Banks (GSIBs). With 3,500+ FS customers across 92 countries, SAS has earned its position as a globally trusted partner.

SAS is also recognised as a leader in risk technology by industry analysts. We are the only vendor to earn a top-five rank in the Chartis RiskTech 100 every year since 2005. We are currently rated number one in the 2023 RiskTech Quadrant for model risk governance solutions, and a category leader in the RiskTech Quadrant for model validation. We also won three prestigious prizes at Risk.net's 2022 Risk Technology Awards.

Our bias detection and mitigation capabilities empower your business teams to address fairness issues quickly and effectively—and we can help you deploy them in weeks. And in the longer term, our cloud-based solution for model lifecycle management and model risk management can scale to thousands of models across geographies, business units, and product segments. We can provide both the robust governance and the agility your firm needs to meet regulatory requirements while accelerating model deployment.

SAS provides a flexible approach based on proven model building, model management, and model risk solutions that are readily available and easy to deploy. By connecting these technologies and integrating them with open-source tools, we can empower your data science teams and provide end-to-end transparency, governance, and auditability throughout the model lifecycle.

Next steps

Reach out to your dedicated SAS account team to learn how we can help your firm hit your Consumer Duty deadlines and build a sustainable framework for meeting all your AI ethics and compliance goals.

sas.com/uk/banking